THE MAXIMUM-LIKELIHOOD DECODING THRESHOLD FOR CYCLE CODES OF GRAPHS

PETER NELSON AND STEFAN H.M. VAN ZWAM

Abstract. For a class $\mathcal{C}$ of binary linear codes, we write $\theta_{\mathcal{C}}\colon (0,1) \to [0, \frac{1}{2}]$ for the *maximum-likelihood decoding threshold function* of $\mathcal{C}$, the function whose value at $R \in (0,1)$ is the largest bit-error rate p that codes in $\mathcal{C}$ can tolerate with a negligible probability of maximum-likelihood decoding error across a binary symmetric channel. We show that, if $\mathcal{C}$ is the class of cycle codes of graphs, then $\theta_{\mathcal{C}}(R) \le \frac{(1-\sqrt{R})^2}{2(1+R)}$ for each R, and show that equality holds only when R is asymptotically achieved by the cycle codes of regular graphs.

1. Introduction

For a class $\mathcal{C}$ of binary linear codes and for some rate $R \in (0,1)$, we consider the *maximum-likelihood decoding threshold* $\theta_{\mathcal{C}}(R)$ for $\mathcal{C}$ at R. This is the unique $\theta \in [0, \frac{1}{2}]$ such that

- for each $p \in (0, \theta)$ and all $\varepsilon > 0$, given a binary symmetric channel of bit-error rate p, there exists a code $C \in \mathcal{C}$ of rate at least R such that the probability of a error in maximum-likelihood decoding on C is at most ε, and
- for each $p \in (\theta, \frac{1}{2})$ there exists $\varepsilon > 0$ such that, given a binary symmetric channel of bit-error rate p, for each code $C \in \mathcal{C}$ of rate at least R the probability of an error in maximum-likelihood decoding on C is at least ε.

The function $\theta_{\mathcal{C}}(R)$ is the *threshold function* for $\mathcal{C}$; it essentially measures the maximum bit-error rate that can be 'tolerated' by rate-R codes in $\mathcal{C}$ with vanishing probability of a decoding error. Our main result proves an upper bound on this function for the class $\mathcal{G}$ of cycle codes of graphs:

Theorem 1.1. *If $\mathcal{G}$ is the class of cycle codes of graphs and $R \in (0,1)$, then $\theta_{\mathcal{G}}(R) \le \frac{(1-\sqrt{R})^2}{2(1+R)}$. If equality holds, then $R = 1-\frac{2}{d}$ for some $d \in \mathbb{Z}$.*

This research was partially supported by a grant from the Office of Naval Research [N00014-12-1-0031] and by grants from the National Science Foundation Division of Mathematical Sciences [1501985, 1500343].

This generalises a result of Decreusefond and Zémor [4], who proved the same upper bound for the class of cycle codes of regular graphs. Our proof follows theirs conceptually, although our exposition and notation are somewhat different. The proof in [4] implicitly involves a problem of enumerating 'non-backtracking' walks that is trivial for regular graphs but not in general; much of the original material in our proof is related to this difficulty.

When $R = 1 - \frac{2}{d}$ for some $d \in \mathbb{Z}$ (that is, when the cycle codes of large d-regular graphs have rate close to R) our theorem does not improve the bound $\theta_{\mathcal{G}}(R) \leq \frac{(1-\sqrt{R})^2}{2(1+R)}$. In this case, however, the bound is known to be best-possible; Zémor and Tillich [13] showed, when $d-1$ is one of various prime powers, that certain families of d-regular Ramanujan graphs have cycle codes attaining this threshold (that is, can tolerate a bit-error rate of p for any $p < \frac{(1-\sqrt{R})^2}{2(1+R)}$), and later random constructions due to Alon and Bachmat [1] can be demonstrated to give the same result for all $d \geq 3$. Combining these constructions with Theorem 1.1, we have the following:

Theorem 1.2. *If $\mathcal{G}$ is the class of cycle codes of graphs, and $R = 1 - \frac{2}{d}$ for some integer $d \geq 3$, then $\theta_{\mathcal{G}}(R) = \frac{(1-\sqrt{R})^2}{2(1+R)}$.*

Theorem 1.1 implies that this equality holds for no other $R \in (0,1)$; this can be interpreted as a statement that the cycle codes of regular graphs are 'best' among all cycle codes.

Theorem 1.1 will be derived as a consequence of a stronger upper bound for $\theta_{\mathcal{G}}$, given in Theorem 3.4. While the bound in Theorem 3.4 is highly technical in its statement, we believe (Conjecture 3.5) that it is in fact the correct upper bound.

Minor-Closed Classes. The main result of [11] shows that the failure of the cycle codes to be 'asymptotically good' extends to every *proper minor-closed* subclass of binary codes; that is, every proper subclass that is closed under puncturing and shortening. The proof uses a deep result in matroid structure theory due to Geelen, Gerards and Whittle [6] that states, roughly, that the 'highly connected' members of any such class of codes are close to being either cycle codes or their duals.

We believe that this paradigm that the members of any minor-closed subclass of binary codes are 'nearly' cycle or cocycle codes will also apply to the threshold function. We predict that the threshold function $\theta_{\mathcal{G}}(R)$ for any minor-closed class agrees with that of either the class $\mathcal{G}$ of cycle codes or the class $\mathcal{G}^*$ of cocycle codes. It is easily shown (see

[6]) that $\theta_{\mathcal{G}^*}(R) = 0$ for all $R \in (0,1)$. Geelen, Gerards and Whittle [6] made the following striking conjecture:

Conjecture 1.3. *Let $\mathcal{C}$ be a proper subclass of the binary linear codes that is closed under puncturing and shortening. Either*

- *$\mathcal{G} \subseteq \mathcal{C}$ and $\theta_{\mathcal{C}} = \theta_{\mathcal{G}}$, or*
- *$\theta_{\mathcal{C}} = 0$.*

In other words, the presence or absence of the class of cycle codes should be all that determines the threshold function for any minor-closed class. Proving this conjecture would likely require a combination of the matroidal techniques in [11] and the algebraic and probabilistic ideas in this paper.

2. Preliminaries

We give some basic definitions in coding theory that, together with the definition of threshold function in the introduction, are all that are required for this paper; a more comprehensive reference is found in [10]. We also use some standard graph theory terminology from [5] and [7].

For integers $n \geq k \geq 0$, a *binary linear $[n,k]$-code* is a k-dimensional subspace C of some n-dimensional vector space V over GF(2). We call the elements of C *codewords*. The *rate* of C is the ratio $R = \frac{k}{n}$.

2.1. Cycle codes.

This paper is concerned solely with the cycle codes of graphs. For a finite graph $G = (V,E)$, the *cycle code* of G is the subspace of $\mathrm{GF}(2)^E$ whose elements are exactly the characteristic vectors of cycles of G (that is, edge-disjoint unions of circuits of G, or equivalently edge-sets of even subgraphs of G). We write $\mathcal{G}$ for the class of all such codes; it is well-known that every cycle code is the cycle code of a connected graph.

If G is connected, then its cycle code C is a binary linear $[n,k]$-code, where $n = |E|$ and $k = |E| - |V| + 1$, giving $R = 1 - \frac{|V|}{|E|} + \frac{1}{|E|}$. The ratio $\frac{|V|}{|E|}$ is exactly $\frac{2}{\mu(G)}$, where $\mu(G)$ denotes the average degree of G; we adopt this notation $\mu(G)$ throughout the paper. The above formula implies that a large connected graph G has a cycle code of rate $R \approx 1 - \frac{2}{\mu(G)}$. A simple 'error-tolerance' parameter of C is the minimum Hamming distance d between two codewords of C; this is equal to the *girth* of G (the length of a shortest circuit of G) – we will write $d(G)$ for the girth of a graph G.

2.2. Maximum-likelihood decoding.

Suppose that some codeword c of a linear $[n,k]$-code $C \subseteq V$ is transmitted across a binary symmetric channel with bit-error rate $p \in (0, \frac{1}{2})$, giving some $x \in V$ obtained by switching the value of each entry of c independently with probability p. *Maximum-likelihood decoding* (abbreviated ML-decoding) is the process where, given x, we attempt to recover c by choosing the codeword $c' \in C$ with the highest probability to have been sent, given that x has been received. If this choice is ambiguous (that is, if this maximum is not unique) or gives an incorrect answer (that is, if $c' \neq c$), then we say a *decoding error* has been made; this occurs with some probability depending on p and C but, by linearity, not on the particular codeword c. In this particular setting of a constant bit-error probability $p < \frac{1}{2}$ that behaves independently on each bit, ML-decoding is equivalent to *nearest-neighbour* decoding, where c' is simply chosen to be the closest codeword to x in Hamming distance. We remark that our definition of ML-decoding deviates slightly from the standard one, in which a decoding error is also avoided with nonzero probability in the case of an ambiguous choice. This difference will not affect the asymptotic analysis with which we are concerned.

ML-decoding is hard for general binary codes [3], but an attractive property of cycle codes of graphs (and an important motivating factor for this paper) is that ML-decoding can be implemented efficiently for cycle codes using standard techniques in combinatorial optimization (see [12]). This is the case because the probability of a decoding error can be understood purely graphically: if C is the cycle code of a graph $G = (V, E)$ and codewords of C are transmitted across a channel of bit-error rate $p \in (0, \frac{1}{2})$, then the probability of an ML-decoding error is exactly the probability, given a set $X \subseteq E$ formed by choosing each edge uniformly at random with probability p, that X contains at least half of the edges of some circuit of G. Thus, to prove our main theorem, we study random subsets of edges of a graph. From this point on, given a set E and some $p \in [0, 1]$, we refer to a random set $X \subseteq E$ formed by including each element of E independently at random with probability p as a *p-random subset of E*.

3. Non-backtracking walks

A *non-backtracking walk* of length ℓ in a graph G is a walk $(v_0, v_1, \dots, v_\ell)$ of G so that $v_{i+1} \neq v_{i-1}$ for all $i \in \{1, \dots, \ell - 1\}$. In all nontrivial cases, the number of such walks grows roughly exponentially in ℓ; in this section we estimate the base of this exponent, mostly following ([2], Theorem 1).

Let $G = (V, E)$ be a simple connected graph of minimum degree at least 2. Let $\bar{E} = \{(u, v) \in V^2 : u \sim_G v\}$ be the $2|E|$-element set of arcs of G. Let $B = B(G) \in \{0,1\}^{\bar{E}\times\bar{E}}$ be the matrix so that $B_{(u,v),(u',v')} = 1$ if and only if $u' = v$ and $u \neq v'$. It is easy to see that

(1) B is the adjacency matrix of a strongly connected digraph (essentially the 'line digraph' of G), and
(2) For each integer $\ell \geq 1$, the entry $(B^\ell)_{e,f}$ is the number of non-backtracking walks of length $\ell+1$ in G with first arc $(v_0, v_1) = e$ and last arc $(v_\ell, v_{\ell+1}) = f$.

By (1) and the Perron-Frobenius theorem (see [7], section 8.8), there is a positive real eigenvalue λ_* of B and an associated positive real eigenvector w_*, so that $|\lambda_*| \geq |\lambda|$ for every eigenvalue λ of B. Furthermore, by Gelfand's formula [8] we have $\lambda_* = \lim_{n\to\infty} \|B^n\|^{1/n}$, where $\|B^n\|$ denotes the sum of the absolute values of the entries of B^n. By (2), the parameter $\lambda_* = \lambda_*(B(G))$ thus governs the growth of non-backtracking walks in G.

Note that B^ℓ has only nonnegative entries, so $\left\|B^\ell\right\| = \overline{\mathbf{1}}^T B^\ell \overline{\mathbf{1}}$. Let $\mu = \mu(G) = \frac{1}{n}|\bar{E}|$ denote the average degree of G. The proof of Theorem 1 of [2] contains the following:

Lemma 3.1. *Let G be a connected graph of minimum degree at least 2 and let $B = B(G)$. Then $\overline{\mathbf{1}}^T B^\ell \overline{\mathbf{1}} \geq (n\mu)\Lambda^\ell$, where*

$$\Lambda = \Lambda(G) = \prod_{v\in V} (d_G(v) - 1)^{d_G(v)/(n\mu)}.$$

It follows in turn from this lemma that $\left\|B^\ell\right\|^{1/\ell} \geq \Lambda(G)$, so $\lambda_*(B(G)) \geq \Lambda(G)$. As observed in [2], the log-convexity of the function $(x-1)^x$ (for $x > 1$) implies that $\Lambda(G) \geq \mu(G) - 1$. For each $x \in \mathbb{R}$, let $\eta(x) = \min(x - \lfloor x \rfloor, \lceil x \rceil - x)$ denote the distance from x to the nearest integer. The following lemma, which is proved by slightly improving the bound $\Lambda(G) \geq \mu(G) - 1$ when $\mu(G)$ is not an integer, is an unilluminating exercise in calculus.

Lemma 3.2. *Let $\mu_0 \in \mathbb{R}$ satisfy $\mu_0 \geq 2$ and let G be a connected graph with minimum degree at least 2 and average degree at least μ_0. Then $\lambda_*(B(G)) \geq \mu_0 - 1 + \frac{\eta(\mu_0)^3}{8\mu_0^3}$.*

Proof. Let $n = |V(G)|$, let $d_1, \dots, d_n$ be the degrees of the vertices of G, and let $\mu = \frac{1}{n}\sum_{i=1}^n d_i \geq \mu_0$ be the average degree of G. Let $\eta = \eta(\mu)$; note that $\mu \geq 2 + \eta$. Define $g\colon (1,\infty) \to \mathbb{R}$ by $g(x) = x\ln(x-1)$; observe that $g'(x) = \frac{x}{x-1} + \ln(x-1)$ and $g''(x) = \frac{x-2}{(x-1)^2}$. We have $\ln(\Lambda(G)) = \frac{1}{n\mu}\sum_{i=1}^n g(d_i)$; for each i, Taylor's theorem gives

$$g(d_i) = g(\mu) + g'(\mu)(d_i - \mu) + \tfrac{1}{2}g''(\xi_i)(d_i - \mu)^2$$

for some ξ_i between d_i and μ_0. We now estimate the 'error' terms.

Claim 3.2.1. *$\frac{1}{2}g''(\xi_i)(d_i-\mu)^2 \geq \frac{\eta^3}{8\mu^2}$ for each i.*

Proof of claim: First suppose that $d_i = 2$. Then $g(d_i) = 0$, so

$$\begin{aligned}\tfrac{1}{2}g''(\xi_i)(2-\mu)^2 &= -g(\mu) - g'(\mu)(2-\mu)\\ &= (\mu-2)\left(\tfrac{\mu}{\mu-1} + \ln(\mu-1)\right) - \mu\ln(\mu-1)\\ &= \tfrac{\mu(\mu-2)}{\mu-1} - 2\ln(\mu-1).\end{aligned}$$

Note that the above expression is equal to $1.174\ldots > 1$ for $\mu = \frac{7}{3}$, and is increasing in μ for $\mu \in (2,\infty)$. If $\mu \geq \frac{7}{3}$ then we therefore have $\frac{1}{2}g''(\xi_i)(2-\mu)^2 > 1$. If $\mu < \frac{7}{3}$ then $\mu = 2+\eta$ and $\eta < \frac{1}{3}$, so

$$\begin{aligned}\tfrac{\mu(\mu-2)}{\mu-1} - 2\ln(\mu-1) &= \tfrac{\eta(2+\eta)}{1+\eta} - 2\ln(1+\eta)\\ &\geq \tfrac{\eta(2+\eta)}{1+\eta} - 2(\eta - \tfrac{1}{2}\eta^2 + \tfrac{1}{3}\eta^3)\\ &= \tfrac{\eta^3}{3(1+\eta)}(1-2\eta)\\ &> \tfrac{1}{12}\eta^3,\end{aligned}$$

where the last inequality uses $\eta < \frac{1}{3}$. Therefore if $d_i = 2$ we have $\frac{1}{2}g''(\xi_i)(d_i-\mu)^2 \geq \min(1, \frac{1}{12}\eta^3) = \frac{1}{12}\eta^3 > \frac{\eta^3}{8\mu^2}$.

Suppose that $d_i \geq 3$. Since ξ_i is between μ and d_i, we have $\xi_i \geq \min(d_i,\mu) \geq \min(3,\mu)$, so $\xi_i - 2 \geq \eta$. Therefore $g''(\xi_i) \geq \frac{\eta}{(\xi_i-1)^2} > \frac{\eta}{\xi_i^2}$. Thus, using $\xi_i \leq \max(\mu, d_i)$, we have

$$\tfrac{1}{2}g''(\xi_i)(d_i-\mu)^2 \geq \frac{\eta(d_i-\mu)^2}{2\xi_i^2} \geq \frac{\eta(d_i-\mu)^2}{2\max(\mu,d_i)^2}.$$

It is easy to show, since $d_i \in \mathbb{Z}$, that $\left|\frac{d_i-\mu}{\max(\mu,d_i)}\right| \geq \frac{\eta}{\mu+\eta} \geq \frac{\eta}{2\mu}$, so $\frac{1}{2}g''(\xi_i)(d_i-\mu)^2 \geq \frac{\eta^3}{8\mu^2}$ and the claim follows. □

Using the claim, we have

$$\begin{aligned}\ln(\Lambda(G)) &= \frac{1}{n\mu}\sum_{i=1}^{n} g(d_i)\\ &= \frac{1}{n\mu}\sum_{i=1}^{n}\left(g(\mu) + g'(\mu)(d_i-\mu) + \tfrac{1}{2}g''(\xi_i)(d_i-\mu)^2\right)\\ &= \frac{1}{n\mu}\left(ng(\mu) + \sum_{i=1}^{n}\tfrac{1}{2}g''(\xi_i)(d_i-\mu)^2\right)\end{aligned}$$

$$\begin{aligned}&\geq \ln(\mu-1)+\frac{1}{n\mu}\left(\frac{n\eta^3}{8\mu^2}\right)\\&=\ln(\mu-1)+\frac{\eta^3}{8\mu^3}.\end{aligned}$$

So $\Lambda(G) \geq (\mu-1)\exp\left(\frac{\eta^3}{8\mu^3}\right) \geq \mu-1+(\mu-1)\left(\frac{\eta^3}{8\mu^3}\right) \geq \mu-1+\frac{\eta^3}{8\mu^3}$. One easily checks that the function $h(y) = y-1+\frac{\eta(y)}{8y^3}$ is strictly increasing on $(2,\infty)$; since $\mu \geq \mu_0$ and $\lambda_*(B(G)) \geq \Lambda(G)$, it follows that $\lambda_*(B(G)) \geq \mu_0 - 1 + \frac{\eta(\mu_0)^3}{8\mu_0^3}$, as required. □

For each $\mu \geq 2$, let $\mathcal{G}_\mu$ denote the class of connected graphs with average degree at least μ and minimum degree at least 2. For every integer $n \geq \mu+1$, let

$$\lambda_*(\mu; n) = \inf\{\lambda_*(B(G))\colon G \in \mathcal{G}_\mu, |V(G)| = n\},$$

noting that this infimum is finite since $K_n \in \mathcal{G}_\mu$ for all $n \geq \mu+1$. Define $\lambda_*\colon [2,\infty) \to \mathbb{R}$ by $\lambda_*(\mu) = \liminf_{n\to\infty} \lambda_*(\mu; n)$. The following is immediate from Lemma 3.2.

Lemma 3.3. $\lambda_*(\mu) \geq \mu - 1$. *If equality holds, then* $\mu \in \mathbb{Z}$.

Having defined the function λ_*, we can now state the more technical main theorem from which Theorem 1.1 will easily follow.

Theorem 3.4. *If* $\mathcal{G}$ *is the class of cycle codes of graphs and* $R \in (0,1)$, *then* $\theta_{\mathcal{G}}(R) \leq \frac{1}{2}\left(1-\sqrt{1-\frac{1}{\lambda^2}}\right)$, *where* $\lambda = \lambda_*\left(\frac{2}{1-R}\right)$.

As mentioned, we believe the above bound is the true value for $\theta_{\mathcal{G}}$.

Conjecture 3.5. *The bound in Theorem 3.4 holds with equality for all* $R \in (0,1)$.

By Theorem 1.2, this conjecture holds when $R = 1-\frac{2}{d}$ for $d \in \mathbb{Z}$.

4. Covering trees

A *locally finite, infinite rooted tree* (hereafter just a *tree*) is a connected acyclic infinite graph Γ of finite maximum degree together with a particular vertex r called the *root.* Adopting some notation of [4] and [9], for $x \in V(\Gamma)$ we write $|x|$ for the distance of x from r, and we write $x \preceq y$ if x is on the path from r to y. We write $x \wedge y$ for the *join* of x and y, the vertex of largest distance from r that is on both the path from r to x and the path from r to y.

The trees we are interested in are 'covering trees' for finite graphs. Let $G = (V, E)$ be a finite graph of minimum degree at least 2 and let $e = (u, v)$ be an arc of G. The *covering tree of G rooted at e* is the tree $\Gamma = \Gamma_e(G)$ where the root is the length-zero walk (u) of G, the other vertices are the non-backtracking walks of G with first arc e and the children of each walk $(u, v, v_2, \dots, v_\ell)$ of length ℓ are its extensions $(u, v, v_2, \dots, v_\ell, v_{\ell+1})$ to nonbacktracking walks of length $\ell + 1$ (ie. where $v_{\ell+1}$ is adjacent to v_ℓ in G and is not equal to $v_{\ell-1}$). Note that the number of vertices of $\Gamma_e(G)$ at distance ℓ from the root is the total number of length-ℓ non-backtracking walks of G with first arc e, which is exactly the sum of the entries of the e-column of $B(G)^{\ell-1}$.

There is a natural homomorphism that associates each walk with its final vertex; if G has large girth, this map preserves much of the local structure of G. To analyse the ubiquity of cycles in a random sample of edges of G, we follow [4] and study a problem of 'fractional percolation' on covering trees, bounding the probability that, given a p-random subset of $E(\Gamma_e(G))$, there is a long path starting at r that is, in a certain sense, dense with edges in the subset.

Let Γ be such a tree, and let $\alpha \in (0, 1)$. Given $X \subseteq E(\Gamma)$, we say that a finite path $(v_0, v_1, \dots, v_n)$ of Γ is *α-adapted* with respect to X if, for each $i \in \{1, \dots, n\}$, the subpath $(v_0, \dots, v_i)$ contains at least αi edges of X. If $t_1, t_2, \dots,$ is a sequence of positive integers and $T_n = \sum_{i=1}^{n} t_i$ is its sequence of partial sums (with $T_0 = 0$), then we say that a path $(x_0, x_1, \dots, x_n)$ of Γ is *(α, t)-adapted* with respect to X if for each $i \in \mathbb{Z}_{>0}$ for which $T_{i+1} < n$, the path $(x_{T_i}, x_{T_i+1}, \dots, x_{T_{i+1}-1})$ is α-adapted, and also the path $(x_{T_j}, x_{T_j+1}, \dots, x_n)$ is α-adapted, where j is minimal so that $T_{j+1} > n$. Note that any initial subpath of an (α, t)-adapted path is (α, t)-adapted.

We will be considering p-random subsets X of $E(\Gamma)$. We first estimate, with an argument used in ([4], Proposition 2), the probability that a given path is α-adapted with respect to X. Henceforth, we denote the 'relative entropy' between α and p by

$$D(\alpha \| p) = \alpha \ln\left(\frac{\alpha}{p}\right) + (1 - \alpha) \ln\left(\frac{1 - \alpha}{1 - p}\right).$$

We remark that [4] defines $D(\alpha\|p)$ as the negative of this formula.

Lemma 4.1. *Let $0 < p < \alpha < 1$. There exists $c > 0$ so that, if $[x_0, x_1, \dots, x_n]$ is a finite path, and X is a p-random subset of the edges of the path, then*

$$\mathbf{P}\left([x_0, \dots, x_n] \text{ is } \alpha\text{-adapted w.r.t. } X\right) \geq cn^{-5/2} \exp(-nD(\alpha\|p)).$$

Proof. We first make a claim that will simplify the estimate.

Claim 4.1.1. *If $|X| \geq \alpha n$, then there exists $\ell \in \{0, \ldots, n-1\}$ such that the path corresponding to the cyclic ordering $[x_\ell, x_{\ell+1}, \ldots, x_n = x_0, x_1, \ldots, x_\ell]$ is α-adapted with respect to X.*

Proof of claim: For each $i \in \mathbb{Z}_n$, let $t_i = 1 - \alpha$ if the edge $x_i x_{i+1}$ is in X, and $t_i = -\alpha$ otherwise. For $0 \leq j \leq j' \leq n$ let $S(j, j') = \sum_{i=j}^{j'-1} t_i$; observe that if $S(j, j') \geq 0$ then the path from x_j to $x_{j'}$ has an α-fraction of its edges in X. In particular, we have $S(0, n) = |X| - \alpha n \geq 0$. Choose $\ell \in \{0, \ldots, n-1\}$ so that $S(0, \ell)$ is minimized. For $\ell \leq h \leq n$ we have $S(\ell, h) = S(0, h) - S(0, \ell) \geq 0$ and for $1 \leq h \leq \ell$ we have $S(\ell, n) + S(0, h) = S(0, n) + (S(0, h) - S(0, \ell)) \geq 0$. It follows from the observation that ℓ satisfies the claim. □

By the above claim and symmetry, the probability that the path $[x_0, \ldots, x_n]$ is α-adapted is at least $\frac{1}{n}\mathbf{P}(|X| \geq \alpha n)$.

It is straightforward to show using $0 < \alpha < 1$ and Stirling's approximation that all sufficiently large n satisfy

$$\lceil \alpha n \rceil! \leq (\alpha n + 1)\lfloor \alpha n \rfloor! \leq \sqrt{2\pi n^3}\left(\tfrac{\alpha n}{e}\right)^{\alpha n}$$

$$(n - \lceil \alpha n \rceil)! \leq \sqrt{2\pi n}\left(\tfrac{(1-\alpha)n}{e}\right)^{(1-\alpha)n},$$

so Stirling's approximation gives $\binom{n}{\lceil \alpha n \rceil} \geq \frac{1}{\sqrt{2\pi} n^{3/2}}\left(\alpha^\alpha (1-\alpha)^{1-\alpha}\right)^{-n}$ for all large n. All large enough n thus satisfy

$$\begin{aligned}
\tfrac{1}{n}\mathbf{P}(|X| \geq \alpha n) &\geq \tfrac{1}{n}\mathbf{P}(|X| = \lceil \alpha n \rceil) \\
&= \frac{1}{n}\binom{n}{\lceil \alpha n \rceil} p^{\lceil \alpha n \rceil}(1-p)^{n - \lceil \alpha n \rceil} \\
&\geq \frac{1}{\sqrt{2\pi} n^{5/2}}\left(\frac{p^\alpha (1-p)^{1-\alpha}}{\alpha^\alpha (1-\alpha)^{1-\alpha}}\right)^n p^{\lceil \alpha n \rceil - \alpha n}(1-p)^{\alpha n - \lceil \alpha n \rceil} \\
&\geq \frac{p}{\sqrt{2\pi} n^{5/2}} \exp(-nD(\alpha \| p));
\end{aligned}$$

since the probability of a path being α-adapted is clearly positive for all n, some $c \in (0, \frac{p}{\sqrt{2\pi}}]$, obtained by taking a minimum over all small n, satisfies the lemma. □

We say a positive integer sequence $t = (t_i : i \geq 1)$ is *slow* if it is nondecreasing and satisfies $\lim_{n\to\infty} t_n = \infty$ and $\lim_{n\to\infty} \frac{t_{n+1}}{\sum_{i=1}^n t_i} = 0$.

The next lemma is the main technical result of this section. It shows that, if t is a slow sequence, G is a graph, and α and p are chosen so that $\exp(D(\alpha \| p))$ is less than the graph invariant $\lambda_*(B(G))$ of the previous section, then there is some arc e_0 of G for which a p-random subset of $E(\Gamma_{e_0}(G))$ will give an arbitrarily long (α, t)-adapted path with probability bounded away from zero. The independence of δ on n and G in this lemma is crucial.

Lemma 4.2. *For all $0 < p < \alpha < 1$, every slow sequence t, and all $\lambda > \exp(D(\alpha \| p))$, there is some $\delta = \delta(t, \lambda, \alpha, p) > 0$ such that, if $n \geq 1$ is an integer and G is a connected graph of minimum degree at least 2 with $\lambda_*(B(G)) \geq \lambda$, then there is an arc e_0 of G so that, given a p-random subset $X \subseteq E(\Gamma_{e_0}(G))$, we have*

$$\mathbf{P}\left(\Gamma_{e_0}(G) \text{ contains an } (\alpha, t)\text{-adapted path of length } n \text{ w.r.t. } X\right) > \delta.$$

Proof. Let $\lambda_* = \lambda_*(B(G))$. Let $t = (t_i : i \geq 1)$ and $T_\ell = \sum_{i=1}^{\ell} t_i$ for each $\ell \geq 0$. Let $\lambda_0 = \exp(D(\alpha \| p))$ and λ_1, λ_2 be real numbers so that $\lambda_0 < \lambda_1 < \lambda_2 < \lambda$. Note that $\lambda_0 > 1$ and $\lambda_* \geq \lambda$.

Let $\Pi(m)$ denote the probability that a path of length m is α-adapted with respect to a p-random subset of its edges, and for each $\ell \geq 0$ let $f(\ell) = \prod_{i=1}^{\ell} \Pi(t_i)^{-1}$ be the reciprocal of the probability that a path of length T_ℓ is (α, t)-adapted. To determine δ, we first estimate f:

Claim 4.2.1. *There exists $M > 0$ such that $f(\ell+1) \leq M\lambda_2^{T_\ell}$ for all ℓ.*

Proof of claim: Let $c > 0$ be given by Lemma 4.1 for p and α. We have

$$\begin{aligned} f(\ell+1) = \prod_{i=1}^{\ell+1} \Pi(t_i)^{-1} &\leq \prod_{i=1}^{\ell+1} \frac{t_i^{5/2}}{c} \exp\left(D(\alpha \| p) \sum_{j=1}^{\ell+1} t_j\right) \\ &= \lambda_0^{T_{\ell+1}} \prod_{i=1}^{\ell+1} \frac{t_i^{5/2}}{c} \\ &= \lambda_1^{T_{\ell+1}} \prod_{i=1}^{\ell+1} \frac{t_i^{5/2}}{c} \left(\frac{\lambda_0}{\lambda_1}\right)^{t_i} \\ &= \lambda_1^{(1+t_{\ell+1}/T_\ell)T_\ell} \prod_{i=1}^{\ell+1} \frac{t_i^{5/2}}{c} \left(\frac{\lambda_0}{\lambda_1}\right)^{t_i}. \end{aligned}$$

Since $\lambda_0 < \lambda_1 < \lambda_2$ and $t_{\ell+1}/T_\ell \to 0$ and $t_\ell \to \infty$, this expression is at most $\lambda_2^{T_\ell}$ for large enough ℓ. The claim follows by taking a maximum over all small ℓ. □

Set $\delta = M^{-1}(\frac{1}{\lambda_2} - \frac{1}{\lambda})$. Let $\bar{E}$ be the set of arcs of G, let $B = B(G)$ and let w_* be the (strictly positive) eigenvector of B for λ_*, normalised to have largest entry 1. Choose $e_0 \in \bar{E}$ such that $w_*(e_0) = 1$. We show that δ and e_0 satisfy the lemma.

For each $e \in \bar{E}$, let b_e be the standard basis vector in $\mathbb{R}^{\bar{E}}$ corresponding to e, and let $N_h(e_0, e) = b_{e_0}^T B^{h-1} b_e$ be the number of non-backtracking walks of length h in G with first arc e_0 and last arc e.

Let $\Gamma = \Gamma_{e_0}(G)$ and r be the root of Γ. Let $\rho\colon V(\Gamma) \setminus \{r\} \to \bar{E}$ be the map assigning each walk to its last arc. Set $\phi(r) = 1$ and, for each vertex $x \neq r$ of Γ, set $\phi(x) = \lambda_*^{1-|x|} w_*(\rho(x))$. Note that $\phi((e_0)) = w_*(e_0) = 1$ and that, for each $x \neq r$ with $\rho(x) = e$, the sum of $\phi(y)$ over the children y of x is

$$\begin{aligned}
&\lambda_*^{1-(|x|+1)} \sum \big(w_*(e')\colon e' \in \bar{E}, B_{e,e'} = 1\big) \\
&= \lambda_*^{-|x|} b_e^T B w_* \\
&= \lambda_*^{1-|x|} w_*(e) = \phi(x).
\end{aligned}$$

(In other words, ϕ is a *unit flow* on Γ.) It follows that for every $h \geq 0$ and all x with $|x| \leq h$, we have $\sum (\phi(y)\colon y \succeq x, |y| = h) = \phi(x)$.

For $X \subseteq E(\Gamma)$, we say that a vertex v of Γ is (α,t)-*reachable* with respect to X if the path of Γ from r to x is (α,t)-adapted with respect to X; let $R(X)$ denote the set of (α,t)-reachable vertices. Fix ℓ so that $T_\ell \geq n$, and define a random variable $Q = Q(X)$ by

$$Q = f(\ell) \sum_{|x| = T_\ell} \phi(x) 1_{R(X)}(x).$$

The $\phi(x)$ sum to 1 over all x with $|x| = T_\ell$, so $\mathbf{E}(Q) = 1$. We now bound the second moment of Q.

Claim 4.2.2. $\mathbf{E}(Q^2) < \delta^{-1}$.

Proof of claim: We have

$$\mathbf{E}(Q^2) = f(\ell)^2 \sum_{|x|=|y|=T_\ell} \phi(x)\phi(y)\mathbf{P}(x, y \in R(X)).$$

For each $z \in V(\Gamma)$, let $k(z)$ be the maximum integer $k \geq 0$ so that $T_k \leq |z|$. There are edge-disjoint paths of lengths $t_1, t_2, \ldots, t_\ell$ and $t_{k(x\wedge y)+2}, t_{k(x\wedge y)+3}, \ldots, t_\ell$ that all must be α-adapted for both x and y to be in $R(X)$ (the first set of paths make up the path from r to x and the second set are contained in the path from $x \wedge y$ to y), so

$$\begin{aligned}
\mathbf{P}(x, y \in R(X)) &\leq \prod_{i=1}^{\ell} \Pi(t_i) \prod_{i=k(x\wedge y)+2}^{\ell} \Pi(t_i) \\
&= f(k(x \wedge y) + 1) f(\ell)^{-2} \\
&\leq M \lambda_2^{T_{k(x\wedge y)}} f(\ell)^{-2} \\
&\leq M \lambda_2^{|x\wedge y|} f(\ell)^{-2},
\end{aligned}$$

where we use the first claim. Using the fact that $|x \wedge y| \geq 1$ whenever $|x| = |y| = T_\ell$, we have

$$\begin{aligned}
\mathbf{E}(Q^2) &\leq M \sum_{|x|,|y|=T_\ell} \phi(x)\phi(y)\lambda_2^{|x\wedge y|} \\
&= M \sum_{1\leq|z|\leq T_\ell} \lambda_2^{|z|} \sum_{\substack{|x|=|y|=T_\ell \\ x\wedge y=z}} \phi(x)\phi(y) \\
&\leq M \sum_{1\leq|z|\leq T_\ell} \lambda_2^{|z|} \left(\sum_{\substack{|x|=T_\ell \\ x\succ z}} \phi(x) \right)^2 \\
&= M \sum_{1\leq|z|\leq T_\ell} \lambda_2^{|z|}\phi(z)^2 \\
&= M \sum_{i=1}^{T_\ell} \lambda_2^i \sum_{|z|=i} \phi(z)^2.
\end{aligned}$$

If $|z| = i \geq 1$, then $w_*(e) \leq 1$ gives

$$\phi(z)^2 = \lambda_*^{2-2i} w_*(\rho(z))^2 \leq \lambda_*^{2-2i} w_*(\rho(z)).$$

For each $e \in \bar{E}$, the number of $z \in V(\Gamma)$ with $|z| = i$ and $\rho(z) = e$ is $N_i(e_0, e) = b_{e_0}^T B^{i-1} b_e$, so since $Bw_* = \lambda_* w_*$ and $w_*(e_0) = 1$, we have

$$\sum_{|z|=i} \phi(z)^2 \leq \lambda_*^{2-2i} b_{e_0}^T B^{i-1} \sum_{e\in\bar{E}} b_e w_*(e) = \lambda_*^{2-2i} b_{e_0}^T B^{i-1} w_* = \lambda_*^{1-i} \leq \lambda^{1-i}.$$

Thus $\mathbf{E}(Q^2) < M \sum_{i=1}^\infty \lambda_2^i \lambda^{1-i} = M(\frac{1}{\lambda_2} - \frac{1}{\lambda})^{-1} = \delta^{-1}$. □

Now by the Cauchy-Schwartz inequality we have

$$1 = \mathbf{E}(Q)^2 = \mathbf{E}(Q \cdot 1_{Q>0})^2 \leq \mathbf{E}(Q^2)\mathbf{E}(1_{Q>0}^2) < \delta^{-1}\mathbf{P}(Q > 0),$$

so $\mathbf{P}(Q > 0) > \delta$. Therefore Γ has an (α, t)-adapted path of length T_ℓ with respect to X with probability greater than δ. Such a path contains an (α, t)-adapted path of length n, giving the result. □

5. Graphs

For a graph $G = (V, E)$ and for $p, \beta \in [0, 1]$, let $f_p^\beta(G)$ denote the probability, given a p-random subset $X \subseteq E$, that X contains at least a β-fraction of the edges of some circuit of G. Recall that $\lambda_*(\mu_0)$ is some value not less than $\mu_0 - 1$.

Theorem 5.1. *For all $\mu_0 \geq 2$ and $0 < p < \beta < 1$ satisfying $\exp(D(\beta\|p)) < \lambda_*(\mu_0)$, there exists $\delta = \delta(\mu_0, p, \beta) > 0$ such that, if G is a connected graph with $\mu(G) \geq \mu_0$, then $f_p^\beta(G) \geq \delta$.*

Proof. It suffices to show this just for graphs of minimum degree at least 2, since deleting a degree-1 vertex from a graph G with $\mu(G) \geq 2$ does not change f_p^β or connectedness, and does not decrease $\mu(G)$. Suppose that the result fails. Then there exists a sequence $G_1, G_2, \dots,$ of graphs of average degree at least μ_0 and minimum degree at least 2, such that $\lim_{n\to\infty}(f_p^\beta(G_n)) = 0$. We clearly have $f_p^\beta(G) \geq p^{d(G)}$ for every graph (this is the probability of a p-random subset containing *every* edge in a given shortest cycle), so we may assume by taking a subsequence that $d(G_i) \geq i$ for each i.

Claim 5.1.1. *There is a slow integer sequence $t = (t_k \colon k \geq 1)$ so that $t_{|V(G_k)|} \leq \sqrt{k}$ for each k.*

Proof of claim: Let $(t_k \colon k \geq 1)$ be a nondecreasing, divergent integer sequence in which the integer $\lfloor\sqrt{r}\rfloor$ occurs at least $|V(G_r)|$ times for each $r \geq 1$. (Such a sequence can be chosen to diverge because each integer is only required to occur finitely often.) By construction we have $t_{|V(G_k)|} \leq \lfloor\sqrt{k}\rfloor$ for each k. Furthermore, if $\ell \geq 1$ and $t_{\ell+1} = d+1 \geq 2$ then the integer d has occured at least $|V(G_{d^2})| \geq d^2$ times before $t_{\ell+1}$, so $t_{\ell+1}/\sum_{i=1}^{\ell} t_i \leq (d+1)/d^3$. It follows that $\lim_{n\to\infty} t_{n+1}/\sum_{i=1}^{n} t_i = 0$, so $(t_k \colon k \geq 1)$ is slow. □

Note that $D(x\|p)$ is increasing in x for $x > p$. Since $\exp(D(\beta\|p)) < \lambda_*(\mu_0)$ we can choose $\alpha \in (\beta, 1)$ and λ' so that

$$\exp(D(\beta\|p)) < \exp(D(\alpha\|p)) < \lambda' < \lambda_*(\mu_0).$$

Let k_0 be large enough so that $\lambda_*(\mu_0; n) \geq \lambda'$ for all $n \geq k_0$. Let $\delta = \delta(t, \lambda', \alpha, p) > 0$ be given by Lemma 4.2. We argue that if k is sufficiently large so that $k \geq k_0$ and $\frac{2\sqrt{k}+1}{k} \leq \alpha - \beta$, then the graph $G = G_k$ satisfies $f_p^\beta(G) \geq \delta$. This contradicts $\lim_{n\to\infty} f_p^\beta(G_n) = 0$.

Let $G = G_k$ for such a k, and let $\Gamma = \Gamma_e(G)$ be the covering tree of G with respect to the arc $e = (r, s)$ given by Lemma 4.2. Let $\pi \colon V(\Gamma) \to V(G)$ assign each path to its final vertex. Since $|V(G)| \geq k \geq k_0$, we have $\lambda_*(B(G)) \geq \lambda_*(\mu_0; k) \geq \lambda'$ by the choice of k_0.

We now relate $f_p^\beta(G)$ to the probability that a p-random subset of $E(\Gamma)$ gives a long (α, t)-adapted path. For each set $Z \subseteq V(G)$, let $G(Z)$ denote the subgraph of G induced by Z.

Recalling notation from the proof of Lemma 4.2, for $X \subseteq E(G)$ we say a vertex v of G is *reachable* with respect to X if $v = r$, or there is an (α, t)-adapted path of G (with respect to X) having first arc e and last vertex v. We write $R(X)$ for the set of all such vertices. Similarly, for $Y \subseteq E(\Gamma)$, we say a vertex v of Γ is *reachable* with respect to Y if there is an (α, t)-adapted path of Γ (with respect to Y) from the root to v. Let $R(Y)$ denote the set of all such vertices. Note, for any X and Y, that each of the sets $R(X)$ and $\pi(R(Y))$ either is equal to $\{r\}$, or induces a connected subgraph of G containing r and s.

Suppose that X is a p-random subset of $E(G)$ and Y is a p-random subset of $E(\Gamma)$. Let C_G denote the event that $G(R(X))$ contains a circuit, and C_Γ denote the event that $G(\pi(R(Y)))$ contains a circuit.

Claim 5.1.2. $\mathbf{P}(C_G) = \mathbf{P}(C_\Gamma)$.

Proof of claim: Let $\mathcal{Z}'$ denote the family of subsets of $V(G)$ that induce an *acyclic* connected subgraph of G containing r and s, and let $\mathcal{Z} = \mathcal{Z}' \cup \{\{r\}\}$. The event C_G fails to hold exactly when $R(X) \in \mathcal{Z}$, so

$$1 - \mathbf{P}(C_G) = \sum_{Z \in \mathcal{Z}} \mathbf{P}(R(X) = Z).$$

Similarly, we have

$$1 - \mathbf{P}(C_\Gamma) = \sum_{Z \in \mathcal{Z}} \mathbf{P}(\pi(R(Y)) = Z).$$

If $Z = \{r\}$, then clearly $\mathbf{P}(R(X) = Z) = \mathbf{P}(\pi(R(Y)) = Z) = 1 - p$. Suppose that $Z \in \mathcal{Z}'$. By acyclicity of $G(Z)$, there is a unique subtree Γ_Z of Γ that contains the root of Γ and satisfies $\pi(V(\Gamma_Z)) = Z$, and moreover $G(Z)$ and Γ_Z are isomorphic finite trees. Now $G(Z)$ and Γ_Z have the same number of edges, and the number of edges of G with exactly one end in $Z \setminus \{r\}$ is equal to the number of edges of Γ with exactly one end in $V(\Gamma_Z)$, so

$$\mathbf{P}(R(X) = Z) = \mathbf{P}(R(Y) = V(\Gamma_Z)) = \mathbf{P}(\pi(R(Y)) = Z).$$

The claim now follows from the two summations above. □

Claim 5.1.3. $\mathbf{P}(C_\Gamma) \geq \delta$.

Proof of claim: By Lemma 4.2, the tree Γ contains, with probability at least δ, a length-$|V(G)|$ path $[v_1, v_2, \dots]$ that is (α, t)-adapted with respect to Y. For any such path, there must be some $i < j$ so that $\pi(v_i) = \pi(v_j)$; now $\{\pi(v_i), \pi(v_{i+1}), \dots, \pi(v_j)\}$ is the vertex set of a closed non-backtracking walk of $G(\pi(R(Y)))$, which must contain a circuit. This implies the claim. □

Claim 5.1.4. $f_p^\beta(G) \geq \mathbf{P}(C_G)$.

Proof of claim: Suppose that $X \subseteq E$ satisfies C_G; i.e. $G(R(X))$ contains a circuit C. It suffices to show that X contains a β-fraction of the edges of some circuit of G. Let $V(C) = [x_0, x_1, \ldots, x_m]$, where x_0 is the end of a shortest (α, t)-adapted path P_0 from r to $V(C)$. If there is some $i \in \{1, \ldots, m\}$ such that there exists in G an (α, t)-adapted path P_i from r to x_i not containing x_{i-1} and an (α, t)-adapted path P_{i-1} from r to x_{i-1} not containing x_i, then $E(P_i) \cup E(P_{i-1}) \cup \{x_{i-1}x_i\}$ contains a circuit C' of G. Moreover, this circuit is the disjoint union of the edge $x_{i-1}x_i$, a set of subpaths that are α-adapted with respect to X, and at most two extra subpaths each of length at most $t_{|V(G)|}$ (these two subpaths are 'partial' subpaths arising because the last intersection point of P_{i-1} and P_i need not cleanly divide these paths into a union of α-dense subpaths), so $|X \cap E(C')| \geq \alpha|E(C')| - 2t_{|V(G)|} - 1$. Now $G = G_k$, so $|E(C')| \geq d(G) \geq k$ and $t_{|V(G)|} \leq \sqrt{k}$, giving

$$\tfrac{|X \cap E(C')|}{|E(C')|} \geq \alpha - \tfrac{2t_{|V(G)|}+1}{|E(C')|} \geq \alpha - \tfrac{2\sqrt{k}+1}{k} \geq \beta,$$

so X contains a β-fraction of the edges of C'.

If no such i exists, then an easy inductive argument implies for each $j \geq 1$ that every (α, t)-adapted path from r to x_j passes through x_{j-1}, so $E(P_0) \cup E(C) - \{x_0 x_m\}$ is the edge set of an (α, t)-adapted path from r to x_m. By a similar argument to the above, we have $|E(C) \cap X| \geq \alpha|E(C)| - 2t_{|V(G)|} - 1$, and thus X contains a β-fraction of the edges of C, giving the claim. □

The last three claims give $f_p^\beta(G) \geq \delta$, implying the theorem. □

6. The Threshold

We now prove Theorems 3.4 and 1.1. Recall that, if C is the cycle code of a graph G, then the probability of a maximum-likelihood decoding error in C over a channel of bit-error rate $p \in (0, \frac{1}{2})$ is exactly the parameter $f_p^{1/2}(G)$ of the previous section. We use this fact to derive Theorem 3.4 (restated here) from Theorem 5.1.

Theorem 6.1. *If $R \in (0,1)$ and $\mathcal{G}$ is the class of cycle codes of graphs, then $\theta_{\mathcal{G}}(R) \leq \frac{1}{2}\left(1 - \sqrt{1 - \frac{1}{\lambda^2}}\right)$, where $\lambda = \lambda_*(\frac{2}{1-R})$.*

Proof. Fix $R \in (0,1)$, let $\mu = \frac{2}{1-R}$ and let $\theta = \frac{1}{2}\left(1 - \sqrt{1 - \frac{1}{\lambda^2}}\right)$, where $\lambda = \lambda_*(\mu)$. Note that $\exp(D(\frac{1}{2}\|\theta)) = \lambda \geq \mu - 1 > 1$ by Lemma 3.3. It is enough to show that for all $p \in (\theta, \frac{1}{2})$ there is some $\varepsilon > 0$ such that the probability of an error in maximum-likelihood decoding of a cycle code of rate at least R, over a binary symmetric channel with bit-error rate p, is at least ε.

Let $p \in (\theta, \frac{1}{2})$. Since $p > \theta$ we have $\exp(D(\frac{1}{2}\|p)) < \lambda$; let $\lambda_0 \in (\exp(D(\frac{1}{2}\|p)), \lambda)$ and let $\mu_0 = \lambda_0 + 1$. Let $\delta = \delta(\mu_0, p, \frac{1}{2})$ be given by Theorem 5.1 and set $\varepsilon = \min(\delta, p^b)$, where $b = \frac{2\mu\mu_0}{\mu - \mu_0}$.

Let C be a cycle code of rate $R(C) \geq R$ and let G be a connected graph whose cycle code is C. Note, since $R > 0$, that G contains a circuit, so $f_p^{1/2}(G) \geq p^{|E(G)|}$. If $\mu(G) \geq \mu_0$ then $f_p^{1/2}(G) \geq \delta \geq \varepsilon$ by Theorem 5.1. Otherwise

$$1 - \tfrac{2}{\mu} = R \leq R(C) = 1 - \tfrac{2}{\mu(G)} + \tfrac{1}{|E(G)|} < 1 - \tfrac{2}{\mu_0} + \tfrac{1}{|E(G)|},$$

so $|E(G)| < \frac{2\mu\mu_0}{\mu - \mu_0} = b$ and thus $f_p^{1/2}(G) \geq p^b \geq \varepsilon$, as required. □

Finally, we restate and prove Theorem 1.1.

Theorem 6.2. *If $\mathcal{G}$ is the class of cycle codes of graphs and $R \in (0,1)$, then $\theta_{\mathcal{G}}(R) \leq \frac{(1-\sqrt{R})^2}{2(1+R)}$. If equality holds, then $R = 1 - \frac{2}{d}$ for some $d \in \mathbb{Z}$.*

Proof. Let $\mu = \frac{2}{1-R}$ and $\lambda = \lambda_*(\mu)$. By Lemma 3.3 we have $\lambda \geq \mu - 1$ with equality if and only if $\mu \in \mathbb{Z}$. Theorem 6.1 thus gives $\theta_{\mathcal{G}}(R) \leq \frac{1}{2}\left(1 - \sqrt{1 + \frac{2}{\mu - 1}}\right)$, with equality only if $\mu \in \mathbb{Z}$: that is, if and only if $R = 1 - \frac{2}{d}$ for some $d \in \mathbb{Z}$. The result now follows from the definition of μ and a computation. □

Acknowledgements. We thank the two anonymous referees for their helpful suggestions that improved the quality of the paper.

References

[1] N. Alon and E. Bachmat, Regular graphs whose subgraphs tend to be acyclic, Random Struct. Algo. 29 (2006), 324–337.

[2] N. Alon, S. Hoory and N. Linial, The Moore Bound for Irregular Graphs, Graph Combinator. 18 (2002), 53–57.

[3] E.R. Berlekamp, R.J. McEliece and H.C.A. van Tilborg, On the inherent intractability of certain coding problems, IEEE Trans. Inform. Theory 24 (1978), 384–386.

[4] L. Decreusefond and G. Zémor, On the error-correcting capabilities of cycle codes of graphs, Combin. Probab. Comput. 6 (1997), 27–38.

[5] R. Diestel, Graph Theory, Springer, 2000.

[6] J. Geelen, B. Gerards and G. Whittle, The highly connected matroids in minor-closed classes, Ann. Comb. 19 (2015), 107–123.

[7] C. Godsil and G. Royle, Algebraic Graph Theory, Springer, 2001.

[8] I. Gelfand, Normierte ringe, Rech. Math. [Mat. Sbornik] N.S., 9 (1941), 3–24

[9] R. Lyons, Random walks and percolation on trees, Ann. Probab. 18 (1990), 931–958.

[10] F. J. MacWilliams and N. J. A. Sloane, The Theory of Error-Correcting Codes, Amsterdam, The Netherlands: North-Holland, 1977.

[11] P. Nelson and Stefan H.M. van Zwam, On the existence of asymptotically good linear codes in minor-closed classes, IEEE Trans. Inform. Theory 61 (2015), 1153–1158.

[12] S.C. Ntafos and S.L. Hakimi, On the complexity of some coding problems, IEEE Trans. Inform. Theory 27 (1981), 794–796.

[13] J-P. Tillich, G. Zémor, Optimal cycle codes constructed from Ramanujan graphs, SIAM J. Discrete Math 10 (1997), 447–459.